The Power of Vocabulary: The Case of Cyclotomic Polynomials *

Jacques Carette Department of Computing and Software
McMaster University,Hamilton, Ontario
carette@mcmaster.ca

James H. Davenport Department of Computer Science
University of Bath, Bath BA2 7AY, United Kingdom
J.H.Davenport@bath.ac.uk

arXiv:1002.0012v1 [cs.SC] 29 Jan 2010

ABSTRACT

We observe that the vocabulary used to construct the "answer" to problems in computer algebra can have a dramatic effect on the computational complexity of solving that problem. We recall a formalization of this observation and explain the classic example of sparse polynomial arithmetic. For this case, we show that it is possible to extend the vocabulary so as reap the benefits of conciseness whilst avoiding the obvious pitfall of repeating the problem statement as the "solution".

It is possible to extend the vocabulary *either* by irreducible cyclotomics *or* by $x^n - 1$: we look at the options and suggest that the pragmatist might opt for both.

1. INTRODUCTION

While sparse polynomials are a natural data structure for human beings (who writes $x^{10}+0x^9+0x^8+0x^7+0x^6+0x^5+0x^4+0x^3+0x^2+0x^1-1$?) and computer algebra systems, algorithms to do more than add and multiply are scarce on the ground, and most texts slip silently from considering sparse polynomials to considering dense ones [8]. This is partly because of the existence of examples showing that the output can be exponentially larger than the input, and hence "nothing can be done". We contend that these examples are basically all cases of the cyclotomic polynomials in disguise, and that, by admitting these to the *output* language, as Schinzel's K-operator [18] effectively does, these examples cease to be absolute barriers to efficient algorithms. Cyclotomic factors can often be recognised relatively efficiently [6], though the worst-case is NP-hard from this result.

THEOREM 1 ([16, THEOREM 6.1]). *It is NP-hard to solve the problem, given a polynomial $p(x) \in \mathbb{Z}[x]$, to determine if p has a root r of modulus 1.*

This is a paradigmatic example of a more general thesis: solving problems in computer algebra requires the concurrent design of the most appropriate vocabulary and algorithms which are polynomial *in the size of the output* so encoded. Naturally, unconstrained multiplication of new vocabulary is not a viable solution, and a methodology for costing this was proposed in [7].

A common problem in computer algebra is "factorize this polynomial". The algorithms commonly used first compute factorizations p-adically, and then deduce the "true" factorization over $\mathbb{Z}$. The traditional approaches [23] are theoretically exponential in the number of p-adic factors, though in practice the exponential aspect can be "controlled" [1]. Polynomial-time (in the degree, and *a fortiori* in the number of p-adic factors) algorithms are known [12], but in practice tend to be slower. The most recent progress is in [20], whose algorithm is faster in practice, and the deduction phase is heuristically polynomial time in the number of p-adic factors.

Of course, for a sparse polynomial such as $x^n - 1$, the size of the polynomial is $O(\log n)$, and so an algorithm polynomial in n is still exponential in the size of the input. Are there algorithms which are polynomial in the size of the input? If the output represents the factors as expanded polynomials, this is impossible. There is however a conjecture that the only cases which cause exponential blowups are cyclotomic factors – we will return to this later.

NOTATION 1. *We define the following for a polynomial*

$$f = \sum_{i=0}^{n} a_i x^i = a_n \prod_{i=1}^{n} (x - \alpha_i) :$$

$\#f$ the number of non-zero terms in f, $|\{i : a_i \neq 0\}|$;

$f_{-1} = x^n f(1/x) = \sum_{i=0}^{n} a_{n-i} x^i$;

$f_{-} = f(-x) = \sum_{i=0}^{n} a_i (-x)^i$;

f_e the even part of f, $\sum_{i=0}^{n/2} a_{2i} x^i$;

f_o the odd part of f, $\sum_{i=0}^{n/2} a_{2i+1} x^i$;

f_2 the root-square, or Graeffe[1], of f, $\pm a_n^2 \prod_{i=1}^{n} (x - \alpha_i^2) = f_e^2 - x f_o^2$;

$||f||_\infty = \max_{i=0}^{n} |a_i|$;

*Research begun while both authors were on sabbatical at David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario

[1]There are various conventions in the literature as to how one handles the $\pm$ arising from the parity of n.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...$10.00.

$||f||_2 = \sqrt{\sum_{i=0}^{n} a_i^2}$;

$||f||_1 = \sum_{i=0}^{n} |a_i|$;

$M(f)$ (the Mahler measure of f) $= |a_n| \prod_{i=1}^{n} \max(1, |\alpha_i|)$.

2. CYCLOTOMIC POLYNOMIALS

NOTATION 2. *Let $d(n)$ denote the number of divisors of the number n (including 1 and n itself).*

THEOREM 2. *It is known [22] that*

$$d(n) \leq n^{(\log 2 + o(1))/\log\log n}. \quad (1)$$

However, we should note the caveats about the distribution of $d(n)$ given at [11, Theorem 432], in particular that it has *average* order $\log n$ but *normal* order roughly $(\log n)^{\log 2}$.

DEFINITION 1. *We will say that a polynomial is* cyclotomic*, if all its roots are roots of unity. Many authors reserve this for irreducible polynomials, but we will explicitly say "irreducible" when we need to.*

NOTATION 3. *Let Φ_k be the k-th irreducible cyclotomic polynomial:*

$$\Phi_k(x) = \prod_{\gcd(j,k)=1} (x - e^{2\pi i j/k}). \quad (2)$$

We denote by C_n the cyclotomic polynomial with all n-th roots of unity, i.e. $C_n = x^n - 1$.

We should note that it is *not* the case that the coefficients of Φ_k are 0 or ± 1. The first counterexample is Φ_{105}, which contains the terms $-2x^7$ and $-2x^{41}$. Φ_{385} contains the terms $-3x^{120}$ and $-3x^{121}$ The growth rate is in fact greater than one might expect, and Φ_{15015} has terms of $23\,x^{2294}$ and $23\,x^{3466}$. $15015 = 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13$, and this looks like the recipe (confirmed in [4]) to make large values[2], but in fact 23 is first attained at $11305 = 5 \cdot 7 \cdot 17 \cdot 19$, as shown in table 1. We note the spectacular leap at $40755 = 3 \cdot 5 \cdot 11 \cdot 13 \cdot 19$, which is the largest coefficient up to $k = 80,000$. [2, Table 3] shows more such leaps for (much) larger n.

THEOREM 3. *[21, Theorem 1] shows that, for infinitely many n,*

$$\log ||\Phi_n||_\infty > \exp\left(\frac{(\log 2)(\log n)}{\log\log n}\right), \quad (3)$$

and indeed this is precisely the right order of (worst-case) growth [3], perhaps better expressed as

$$\limsup_{n\to\infty} \frac{||\Phi_n||_\infty}{\log n/\log\log n} = \log 2. \quad (4)$$

PROPOSITION 4. *$x^n - 1 = \prod_{d|n} \Phi_d(x)$, and these factors are irreducible.*

PROPOSITION 5. *$\Phi_n(x) = \prod_{d|n} (x^d - 1)^{\mu(n/d)}$, where μ is the Möbius function.*

PROPOSITION 6. *$x^n - 1$ has $d(n)$ irreducible factors.*

[2]It *does* give us (exactly) 500 at $255255 = 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 \cdot 17$.

Cyclotomic polynomials are the bugbear of anyone who tries to deal with sparse polynomials.

EXAMPLE 1. *Asking for the factorization, or even the degrees of the factors, of $x^n - 1$ is tantamount to factoring n, since for every prime p dividing n, there is an $x^{p-1} + \cdots$ in the factorization of $x^n - 1$.*

EXAMPLE 2. *Similarly, asking for the degree of Φ_k is, if k is $p \cdot q$ (p, q distinct primes), tantamount to factoring k, since $\phi(k) = (p-1)(q-1)$ and so*

$$p, q = \frac{1}{2}\left(k + 1 - \phi(k) \pm \sqrt{k^2 - 2k - 2k\phi(k) + (\phi(k) - 1)^2}\right).$$

Cyclotomic polynomials are frequently used as examples.

EXAMPLE 3. *[20, p. 185] gives this example*

$$x^{128} - x^{112} + x^{80} - x^{64} + x^{48} - x^{16} + 1,$$

and states that his algorithm sped up Maple by a factor of 500 on this example. From a cyclotomic-aware point of view, such as [6], this polynomial is easy. Four applications of Graeffe's root-squaring process show (as is obvious to the eye) that this is $f(x^{16})$ where

$$f(x) = x^8 - x^7 + x^5 - x^4 + x^3 - x + 1.$$

Another application takes f to itself, and hence f, and so the original polynomial, is cyclotomic. If α is a root of f, $\alpha^{15} = 1$, so $f = \Phi_{15}$, and the original polynomial is

$$\Phi_{15}\Phi_{30}\Phi_{60}\Phi_{120}\Phi_{240}.$$

EXAMPLE 4. *If p is prime, $f = x^p - 1$, then $\#f = 2$ but $f = (x-1)(x^{p-1} + \cdots + 1)$: two factors with 2 and p terms respectively.*

EXAMPLE 5. *If p, q are distinct primes, $f = (x^p - 1)(x^q - 1)$, then $\#f = 4$ but $f = (x-1)^2(x^{p-1} + \cdots + 1)(x^{q-1} + \cdots + 1)$. The square-free decomposition of f is therefore a repeated factor with 2 terms and a factor g with $pq - p - q + 2$ terms respectively. The largest coefficient in g is $\min\{p, q\}$, and $||g||_1 = pq$.*

Obviously, a square-free decomposition of f was a bad idea in this case: however previously-proposed algorithms, e.g. [5, p. 69] tend to do this.

It could be argued that the problem in this case is the 'cofactor', but life is not that simple.

EXAMPLE 6. *If p, q are distinct primes, $f = (x^p - 1)^2(x^q - 1)$, then $\#f = 6$ but the square-free factorization is*

$$(x-1)^3\left(x^{p-1} + \cdots + 1\right)^2\left(x^{q-1} + \cdots + 1\right),$$

and we are forced to write out the large squared factor. The largest coefficient of $\left(x^{p-1} + \cdots + 1\right)^2$ is p, so we had also better not compute it and then take its square root.

It is the contention of this paper that all these difficulties *except the first* are caused by an inadequate vocabulary: the first seems to be intrinsic, in the fact the factorization of numbers can be encoded as a problem of factorization of polynomials. All we can do is recognise the fact.

Table 1: Large coefficients in Φ_k

$\lvert a_i\rvert$	2	3	4	5	6	7	8=9	14	23	25	27	59	359
first Φ_k	105	385	1365	1785	2805	3135	6545	10465	11305	17225	20615	26565	40755
$\phi(k)$	48	240	576	768	1280	1440	3840	6336	6912	10752	12960	10560	17280

[A larger version, independently computed, is in [2].]

3. REPRESENTATIONAL COMPLEXITY

Information theory, whether through the guise of Kolmogorov Complexity [13] or Minimum Description Length [10], tell us that *good* representations of structured objects are two-part codes: a model and an encoding of data using that model. In other words, the proper "length" of an object consists in counting the length of the representation of the model as well as the representation of the data encoded using this model.

In [7], these results from information theory are rephrased so as to apply more directly to Computer Algebra Systems, and applications to simplification are outlined. The basic result is that for large enough structured expressions, it is *always* worthwhile to first formalize the "structure", and then encode the data in such a way as to abstract out that structure. Note that, if model extensions are not allowed, then simplification reduces simply to length reduction. It is exactly the confusion between issues of the (background) model-class and its use in model reduction which caused Moses [15] to argue that "simplification" was impossible to formalize.

When tackling a particular situation, [7] boils down to finding the right *vocabulary* in which to express ones' result. In some cases, the right vocabulary is somewhat counter-intuitive. For example, in the case of algebraic numbers, it was long ago discovered that using minimal polynomials to "encode" an algebraic number was best – although this can seem puzzling in the setting of "solving" a polynomial, as then the answer to the problem is just an encoding of the question. We will return to this issue later.

For the particular case of factoring of polynomials, what does this tell us? All the theoretical results point in the same direction: cyclotomic polynomials are somehow "special cases", especially when one is factoring sparse polynomials. Conventional wisdom already tells us that both dense and sparse polynomials are useful model classes, and that we should have both at our disposal. The mathematical theory of factoring polynomials (as outlined in the rest of this paper) informs us that cyclotomics are undeniably part of the domain of discourse. Combining these together tells us that they should also be part of our model classes. The only remaining question then is whether adding this particular vocabulary actually leads to a *simplification*. To evaluate this, we need to actually display some data structures designed with this new vocabulary, and then evaluate if we have made any real gains.

4. DATA STRUCTURES

Since we will be arguing on the size of data structures, we will need to define our data representations. The precise details might vary, though in practice the conclusions will not. For definiteness, we describe our choices according to the *unaligned* packed encoding of ASN.1 [19]: note that their `SEQUENCE` is what C programmers would think of as `struct`. Our encodings are intended to be practical, though we ignore issues of alignment to word boundaries, and indeed a number of $\lceil \ldots \rceil$ operators are also omitted.

In theory, one needs to have arbitrary sized data fields, which means one needs fields for the length of the size fields. To avoid this, we will assume that N and K are global parameters for the size of the object, bounding the degrees and the size of the coefficients. Since a polynomial's factors always have smaller degree, we can assume $N = \log_2 n$. It is not so easy for the coefficients [14], but we will assume a *single* field of K bits associated with each outermost data structure, giving the size k of all the coefficients stored in that structure. Since there is one of these, we can ignore its cost.

We next give explicit representations for dense polynomials, sparse polynomials, factored polynomials, Φ-aware factorizations and C-aware factorizations (explained fully below).

4.1 A single dense polynomial

We choose a dense representation with a uniform size bound for all the coefficients. A single dense polynomial of degree n requires $\log_2 n$ bits to represent the degree, and then there are $n+1$ coefficients. Hence if each of them requires k bits, we need $\log_2 k$ bits for the telling us this, and then the coefficients require $k+1$ bits (including sign).

$$(k+1)(n+1) + \log_2 k + \log_2 n \tag{5}$$

In pseudocode[3], this might be represented as follows.

```
DensePoly ::= SEQUENCE {
    Degree INTEGER (0...2^N - 1),
    k      INTEGER (0...2^K - 1),
    Coefficients SEQUENCE {
        INTEGER (-2^k...2^k - 1)
        } SIZE (Degree + 1)
}
```

4.2 A single sparse polynomial

We choose a sparse representation with a uniform size bound for all the coefficients. Furthermore, we assume that there are t non-zero coefficients, i.e. t terms to be represented, and t is bounded by the same bound as the degree. A single sparse term from a polynomial of degree n requires $\log_2 n$ bits to represent the degree. Hence the total space is given by

$$\log_2 n + t(k + 1 + \log_2 n). \tag{6}$$

In pseudocode, this might be represented as

```
SparsePoly ::= SEQUENCE {
    TermCount INTEGER (0...2^N - 1),
```

[3]Essentially ASN.1, except that we allow ourselves to write mathematics, enclosed in boxes, in the pseudocode.

```
Terms SEQUENCE {
    Degree INTEGER (0...2^N − 1),
    Coefficient INTEGER (−2^k ... 2^k − 1)
    } SIZE (TermCount)
}
```

Using a Horner scheme might save a few more bits in the representation of the exponent, but rarely appreciably so.

4.3 Representing Factorizations

We will use the same structure for square-free or complete factorizations: the number of (distinct) factors followed by pairs (multiplicity, factor). In pseudocode, this might be represented as

```
Factorization ::= SEQUENCE {
    FactorCount INTEGER (0...2^N − 1),
    Factors SEQUENCE {
        Multiplicity INTEGER (0...2^N − 1),
        Factor Poly
        } SIZE (FactorCount)
}
```

where *Poly* is one of `DensePoly` or `SparsePoly`. Hence with f factors, the overhead (i.e. the cost over and above that of storing the distinct factors themselves) is

$$(f+1)\log_2 n. \tag{7}$$

4.4 Representing Φ-aware Factorizations

We will use the same structure for square-free or complete factorizations: the number of (distinct) factors followed by pairs (multiplicity, factor). However, we do this twice: once for the factors that are Φ_k, and once for those that are not. The factors that *are* Φ_k are stored as k followed by[4] $\phi(k)$. In pseudocode, this might be represented as

```
Factorization ::= SEQUENCE {
    PhiFactorCount INTEGER (0...2^N − 1),
    PhiFactors SEQUENCE {
        Multiplicity INTEGER (0...2^N − 1),
        k INTEGER (0...2^N − 1),
        Degree INTEGER (0...2^N − 1),
        } SIZE (PhiFactorCount),
    FactorCount INTEGER (0...2^N − 1),
    Factors SEQUENCE {
        Multiplicity INTEGER (0...2^N − 1),
        Factor Poly
        } SIZE (FactorCount)
}
```

where *Poly* is one of `DensePoly` or `SparsePoly`. Hence with l Φ_k factors, the cost of storing them is $(3l+1)\log n$. *Any* Φ_k stored this way is cheaper than in any of the previous representation (dense, sparse or factored), so the worst case is when there are no Φ_k in the factorization, when the overhead is merely the one field `PhiFactorCount`.

[4]The inclusion of $\phi(k)$ avoids the problem in example 2.

4.5 Representing C-aware Factorizations

Instead of storing the Φ_k, we could store the complete cyclotomic polynomials x^k-1. Again, we will use the same structure for square-free or complete factorizations: the number of (distinct) factors followed by pairs (multiplicity, factor). Also in this case, we do this twice: once for the factors that are of the form x^k-1, and once for those that are not. One might think to store the factors that *are* x^k-1 simply as k. However, as pointed out in example 1, this will not allow us to answer questions such as "how many factors" in a reasonable time. Hence we store k followed by its prime factorization.

We should also note that Proposition 5 means that we now need *negative* multiplicities as well. In pseudocode, this might be represented as

```
Factorization ::= SEQUENCE {
    CFactorCount INTEGER (0...2^N − 1),
    CFactors SEQUENCE {
        Multiplicity INTEGER (−2^N ... 2^N − 1),
        Degree INTEGER (0...2^N − 1),
        NumFactors INTEGER (0...2^N − 1),
        KFactor SEQUENCE {
            INTEGER (0...2^N − 1)
            } SIZE (NumFactors),
        } SIZE (CFactorCount),
    FactorCount INTEGER (0...2^N − 1),
    Factors SEQUENCE {
        Multiplicity INTEGER (0...2^N − 1),
        Factor Poly
        } SIZE (FactorCount)
}
```

where *Poly* is one of `DensePoly` or `SparsePoly`. Hence with f factors of x^k-1 involved, the overhead (i.e. the cost over and above that of storing the distinct factors themselves) is

$$(f+1)\log_2 n. \tag{8}$$

An alternative formulation might store the factors of k with multiplicity: there does not seem to be a great deal to choose between them. In either case, we should note that asking for the *number* of irreducible factors is no longer trivial, since the number of irreducible factors corresponding to a single x^k-1 is $d(k)$. Furthermore, since we are allowed negative exponents, representing

$$x^k+1 := \left(x^{2k}-1\right) / \left(x^k-1\right),$$

we have to note that not all the irreducible factors of $x^{2k}-1$ are actually factors of the left-hand side. Nevertheless, since we have stored the prime factorization of k, the problem is efficiently soluble (certainly polynomial time in the size of the representation).

5. REPRESENTING SOME CYCLOTOMIC POLYNOMIALS

We now use each of our representations and compute the size of the results.

Table 2: Representing $x^n - 1 = \prod_{d|n} \Phi_d(x)$

Representation	Fully expanded	Square-free factorization	Factored
Dense	$2(n+1) + \log_2 n$	same as factored	$n^{1+\frac{\log 2}{\log\log n}} \log_2 e$
Sparse	$3\log_2 n$	same as factored	$n^{1+\frac{\log 2}{\log\log n}} \log_2 e$
Φ_k	$3\log_2 n$	same as factored	$(2d(n)+1)\log_2 n$
$x^k - 1$	$3\log_2 n$	same as factored	$\log_2 n$

Table 3: Representing $(x^p-1)(x^q-1) = (x-1)^2\Phi_p(x)\Phi_q(x)$ with $n = p+q$

Representation	Fully expanded	Square-free factorization	Factored
Dense	$2(n+1) + \log_2 n$	$(1+\log_2 n)(n+2) + 4\log_2 n$	$2(n+3) + 6\log_2 n$
Sparse	$4\log_2 n$	$(2n+2)\log_2 n$	$(n+10)\log_2 n$
Φ_k	$4\log_2 n$	$6\log_2 n$	$6\log_2 n$
$x^k - 1$	$4\log_2 n$	$2\log_2 n$	$2\log_2 n$

5.1 Factorization of $x^n - 1$

The sizes of the expanded $x^n - 1$ polynomial in dense, and sparse encodings are obvious. The Φ-aware and C-aware versions are within an additive constant of `DensePoly` / `SparsePoly`. For definiteness, we will use `SparsePoly` based counts.

To understand the size of the factored forms, we need to study their sizes a little more closely. In this case, the factored form and the square-free form coincide. There are $d(n)$ factors, of total degree n, hence $n + d(n)$ terms. By Theorem 3, we can bound[5] the size of the coefficients as $\log_2 e$ times the right-hand side of (3), *viz.* $\frac{\log n}{\log\log n}$, to get the size of a dense polynomial to be

$$(n + d(n))\log_2 e \exp\left(\frac{(\log 2)(\log n)}{\log\log n}\right) + d(n)\log_2 n$$

$$\approx n^{1+\frac{\log 2}{\log\log n}} \log_2 e. \quad (9)$$

In general[6], the factors will be essentially dense, so a sparse encoding will save nothing, but have to pay for the cost of storing the degrees with each coefficient, adding $(n + d(n))\log_2 n$, to give

$$(n+d(n))\left(\log_2 n + \log_2 e \exp\left(\frac{(\log 2)(\log n)}{\log\log n}\right)\right) + d(n)\log_2 n$$

$$\approx n^{1+\frac{\log 2}{\log\log n}} \log_2 e. \quad (10)$$

We should note that the asymptotically dominant term is the coefficient storage in this model, which is contrary to intuition, and even the experimental data in table 1, but this merely shows that the asymptotics will take time to be visible.

The results of this section are summarised in Table 2.

5.2 A square-free factorization

Let use now consider $(x^p - 1)(x^q - 1)$, with p, q distinct primes. The "fully expanded" versions are again obvious. The square-free factorization of $(x^p - 1)(x^q - 1) = (x-1)^2\Phi_p(x)\Phi_q(x)$ involves multiplying out $\Phi_p(x)\Phi_q(x)$. This gives us coefficients of size $O(n)$, in fact $n/2$ assuming that p, q are balanced, taking $(\log_2 n) - 1$ bits to represent the magnitude[7].

In the factored representation, we have three factors, of degrees 1, $p-1$ and $q-1$, i.e. total degree $n-1$. All coefficients are bounded by 1. Hence the total is

$$(n-1+3)(\log_2 n+2)+(3\cdot 2+2)\log_2 n \approx (n+10)\log_2 n. \quad (11)$$

The results of this section are summarised in Table 3.

6. IMPLEMENTATION NOTES

6.1 Cyclotomic-free

It is important to review the encodings of the previous section and notice that for cyclotomic-free cases, these encodings involve constant overhead, independent of the degree and of the number of factors. In fact, by using a bit or two in a header word (which modern computer algebra systems always seem to use in their internal representations), one can choose between these encodings as necessary. In other words, cyclotomic-free polynomials do not have to bear any extra representation cost for this vocabulary extension.

We can also construct various mixed cases, in other words sparse polynomials which factor into a cyclotomic part and a small dense cofactor. The difference in encoding cost is correspondingly mixed, although the end result is similar: adding cyclotomics asymptotically wins.

6.2 Which to choose?

We have posited two encodings for "cyclotomic-aware" representations of factorizations: one in terms of the irreducible cyclotomics Φ_k (section 4.4) and one in terms of the 'complete' cyclotomics $C_k = x^k - 1$ (section 4.5). Tables 2 and 3 make it clear that adding cyclotomics to ones' vocabulary is certainly reprensentionally efficient. But which one should be used? We first summarize some the advantages and disadvantages.

Pro Φ_k: clearly polynomial In the Φ_k-representation, the factorization of $x^{p+q}+x^p+x^q+1$ is $\Phi_2^2\Phi_{2p}\Phi_{2q}$, where as in the C_k-representation it is $C_{2p}C_p^{-1}C_{2q}C_q^{-1}$. Functions meant to extract information from products of polynomials (degrees, multiplicity, etc) still function

[5] In theory, not all the factors *can* have coefficients this large, but the gain from exploiting this is relatively small.

[6] $x^{2^k} - 1$ is an obvious counter-example.

[7] In this case, they are all positive, but we can't count on this in general.

'easily', whereas spotting that the C_k-representation is not square-free is cheap, but not 'obvious'.

Anti Φ_k: worst-case blowup The last two entries of table 2 show that, for highly composite n, representing C_n in terms of Φ_k can require almost n times as much, i.e. exponentially more, since the input has size $O(\log n)$, space, by Theorem 2.

Notes The pragmatist would probably choose Φ. The theoretician would be swayed by the complexity argument and want C. Possibly the best answer is to admit Φ, and C with an additional extension to the vocabulary as in (12).

To further illustrate one particular difficulty, let us consider factoring of C_{105}. It factors into

$$\Phi_1(x)\Phi_3(x)\Phi_5(x)\Phi_7(x)\Phi_{15}(x)\Phi_{21}(x)\Phi_{35}(x)\Phi_{105}(x).$$

But if all we have at our disposal is C, then the best we can do (which is still better than using sparse polynomials) is

$$C_1(x)\frac{C_3(x)}{C_1(x)}\frac{C_5(x)}{C_1(x)}\frac{C_7(x)}{C_1(x)}\frac{C_{15}(x)C_1(x)}{C_3(x)C_5(x)}\frac{C_{21}(x)C_1(x)}{C_3(x)C_7(x)}\times$$

$$\frac{C_{35}(x)C_1(x)}{C_5(x)C_7(x)}\frac{C_{105}(x)C_7(x)C_5(x)C_3(x)}{C_{35}(x)C_{21}(x)C_{15}(x)C_1(x)}.$$

An even more succinct representation is

$$\prod_{\substack{k=1\\ k|105}}^{105}\Phi_k(x) \tag{12}$$

at the cost of another vocabulary extension. This is why in the representation of C in section 4.5, we list the factors of n, which allows us to recover this factorization relatively easily. This still means that in the Φ-aware encoding, obtaining the number of factors, the degrees of each of the factors, or the multiplicity of each of the factors is straightforward, while for the C-aware encoding, these simple questions now require some (small) amount of computation. , and so forth. All of these questions remain just as easy to answer in the Φ-aware encodings as they were before. However, for the C-aware encodings, these "simple" questions now require some actual computations to resolve. In other words, adding Φ to our vocabulary is a very minor change with clear efficiency gains, while adding C is slightly disruptive but with even greater asymptotic efficiency gains.

7. CONCLUSION

In section 5 we have shown how some "troublesome" factorizations, i.e. examples 1 and 5, cease to consume inordinate space when the cyclotomics are represented explicitly. But what of arbitrary polynomials and their factorizations?

The answer is that we do not know, but there are some tantalizing results. [9] state that, provided f is non-reciprocal ($f \neq \pm f_{-1}$), f has a factor with at most $c_2(\#f, ||f||_\infty)$ terms, independent of n, and quotes [17] to say that if the polynomial has no reciprocal factors, then all irreducible factors have at most $c_2(\#f, ||f||_\infty)$ terms. This is a significant step towards controlling the dependence of the *output size* on n, though it falls short of saying that is polynomial in the following ways:

- there is no guarantee that c_2 depends *polynomially* on $\#f$ or $\log ||f||_\infty$, and indeed c_2 is still rather mysterious to the authors [8, Challenge 4];
- nothing is said about the size of the coefficients, and all known technology makes them depend exponentially on n (i.e. the output size depends linearly on n).

We note, though, that the *known* examples of such growth [14, §3], depend on cyclotomic polynomials, so one could hope that the second problem does not occur in practice.

We are convinced, and we hope to have convinced the reader, that, regardless of whether using cyclotomics ultimately makes the problem of sparse polynomial factorization more tractable, even if not guaranteed polynomial (see Theorem 1), in the input size, that they are most definitely worth having in the basic vocabulary used for the *output* of factoring. We strongly recommend that the specification for what it means to factor a polynomial be thus amended.

8. REFERENCES

[1] J.A. Abbott, V. Shoup, and P. Zimmermann. Factorization in $\mathbb{Z}[x]$: The Searching Phase. In C. Traverso, editor, *Proceedings ISSAC 2000*, pages 1–7, 2000.

[2] A. Arnold and M. Monagan. Calculating cyclotomic polynomials of very large height. `http://www.cecm.sfu.ca/~ada26/cyclotomic/CalcCycloPolys.pdf` 2008.

[3] P.T. Bateman. Note on the coefficients of the cyclotomic polynomial. *Bull. AMS*, 55:1180–1181, 1949.

[4] P.T. Bateman, C. Pomerance, and R.C. Vaughan. On the size of the coefficients of the cyclotomic polynomial. *Colloq. Math. Soc. J. Bolyai*, 34:171–202, 1984.

[5] F. Beukers and C.J. Smyth. Cyclotomic points on curves. *Number theory for the millennium*, pages 67–85, 2002.

[6] R.J. Bradford and J.H. Davenport. Effective Tests for Cyclotomic Polynomials. In P. Gianni, editor, *Proceedings ISSAC 1988*, pages 244–251, 1989.

[7] J. Carette. Understanding Expression Simplification. In J. Gutierrez, editor, *Proceedings ISSAC 2004*, pages 72–79, 2004.

[8] J.H. Davenport and J. Carette. The Sparsity Challenges. *To appear in Proc. SYNASC 2009*, 2010.

[9] M. Filaseta, A. Granville, and A. Schinzel. Irreducibility and greatest common divisor algorithms for sparse polynomials. *Number Theory and Polynomials*, pages 155–176, 2008.

[10] Peter D. Grünwald. *The Minimum Description Length Principle*, volume 1 of *MIT Press Books*. The MIT Press, December 2007.

[11] G.H. Hardy and E.M. Wright. An Introduction to the Theory of Numbers (5th. ed.). *Clarendon Press*, 1979.

[12] A.K. Lenstra, H.W. Lenstra, Jun., and L. Lovász. Factoring Polynomials with Rational Coefficients. *Math. Ann.*, 261:515–534, 1982.

[13] Ming Li and Paul Vitanyi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer-Verlag, Berlin, 1997.

[14] M. Mignotte. Some Inequalities About Univariate Polynomials. In *Proceedings SYMSAC 81*, pages 195–199, 1981.

[15] J. Moses. Algebraic Simplification — A Guide for the Perplexed. *Comm. ACM*, 14:527–537, 1971.

[16] D.A. Plaisted. Sparse Complex Polynomials and Irreducibility. *J. Comp. Syst. Sci.*, 14:210–221, 1977.

[17] A. Schinzel. Reducibility of Lacunary Polynomials, I. *Acta Arith.*, 16:123–159, 1969.

[18] A. Schinzel. Selected Topics on Polynomials. *University of Michigan Press*, 1982.

[19] International Telecommunications Union. Information technology — ASN.1 encoding rules: Specification of Packed Encoding Rules (PER). *Standard X.691*, 2002.

[20] M. van Hoeij. Factoring polynomials and the knapsack problem. *J. Number Theory*, 95:167–189, 2002.

[21] R.C. Vaughan. Bounds for the Coefficients of Cyclotomic Polynomials. *Michigan Math. J.*, 21:289–295, 1974.

[22] S. Wigert. Sur l'ordre de grandeur du nombre des diviseurs d'un entier. *Arkiv för Mat. Astr. Fys.*, 3:1–9, 1907.

[23] H. Zassenhaus. On Hensel Factorization I. *J. Number Theory*, 1:291–311, 1969.